

ABSTRACT OF THE DISCLOSURE

A method for concurrent data migration includes classifying files to be migrated into plural jobs, selecting media to which to migrate each job, and using plural drives concurrently to write the jobs to the media. The selection of a medium is performed in a way that prevents the number of writeable media from exceeding the number of available drives, unless no allocated medium has sufficient space to store any files in a migration job. A medium is preferentially selected that has already been allocated for writing, has space to store at least one file in the job, is not in use for another job, and can be robotically mounted on a drive. If such a medium does not exist, then the set of available media is canvassed to locate an alternative medium. The attributes of each medium are evaluated to determine which medium can be selected most consistently with the goals of (1) preventing the number of media from exceeding the number of drives, and (2) providing sufficient media to allow plural drives to be used concurrently. The technique can be embodied in a file management environment that transparently migrates files meeting certain criteria and stores the location of the migrated file in a reparse point provided by the file system.